

Alibaba Cloud Apsara Stack Enterprise

Glossary

Version: 1808..

Issue: 20180831

Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.









1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.
2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company, or individual in any form or by any means without the prior written consent of Alibaba Cloud.
3. The content of this document may be changed due to product version upgrades, adjustments, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and the updated versions of this document will be occasionally released through Alibaba Cloud-authorized channels. You shall pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.
4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides the document in the context that Alibaba Cloud products and services are provided on an "as is", "with all faults" and "as available" basis. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not bear any liability for any errors or financial losses incurred by any organizations, companies, or individuals arising from their download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, bear responsibility for any indirect, consequential, exemplary, incidental, special, or punitive damages, including lost profits arising from the use or trust in this document, even if Alibaba Cloud has been notified of the possibility of such a loss.
5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified,

reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.

- 6.** Please contact Alibaba Cloud directly if you discover any errors in this document.

Generic conventions

Table -1: Style conventions

Style	Description	Example
	This warning information indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	 Danger: Resetting will result in the loss of user configuration data.
	This warning information indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	 Warning: Restarting will cause business interruption. About 10 minutes are required to restore business.
	This indicates warning information, supplementary instructions, and other content that the user must understand.	 Note: Take the necessary precautions to save exported data containing sensitive information.
	This indicates supplemental instructions, best practices, tips, and other contents.	 Note: You can use Ctrl + A to select all files.
>	Multi-level menu cascade.	Settings > Network > Set network type
Bold	It is used for buttons, menus, page names, and other UI elements.	Click OK .
Courier font	It is used for commands.	Run the <code>cd /d C:/windows</code> command to enter the Windows system folder.
<i>Italics</i>	It is used for parameters and variables.	<code>bae log list --instanceid Instance_ID</code>
[] or [a b]	It indicates that it is a optional value, and only one item can be selected.	<code>ipconfig [-all -t]</code>
{ } or {a b}	It indicates that it is a required value, and only one item can be selected.	<code>switch {stand slave}</code>

Contents

Legal disclaimer.....	I
Generic conventions.....	I
Glossary.....	1
A.....	1
B.....	4
C.....	5
D.....	8
E.....	12
F.....	13
G.....	14
H.....	15
I.....	16
J.....	18
K.....	18
L.....	19
M.....	22
N.....	24
O.....	25
P.....	26
R.....	29
S.....	32
T.....	37
U.....	39
V.....	39
W.....	41

Glossary

[A](#) | [B](#) | [C](#) | [D](#) | [E](#) | [F](#) | [G](#) | [H](#) | [I](#) | [J](#) | [K](#) | [L](#) | [M](#) | [N](#) | [O](#) | [P](#) | [R](#) | [S](#) | [T](#) | [U](#) | [V](#) | [W](#) |

A

Additional read throughput

The part of read throughput actually consumed that exceeds the reserved read throughput, and is measured per second.

Administrator

The most advanced user who has the highest administration permissions to systems or businesses.

Alarm contact

The recipient of an alarm notification. Alarm notifications are based on alarm contacts and alarm groups. Alarm notifications can only be received after you create alarm contacts and alarm groups and then select the corresponding alarm group when creating the alarm rule.

Alibaba Cloud Security

A one-stop security service established based on the data analysis capability of the Alibaba Cloud cloud computing platform. It provides security vulnerability detection and webpage Trojan detection for small- and medium-sized websites, and host intrusion detection and anti-DDoS for cloud server users.

Alibaba Cloud Security Manager

An all-around security technology and consultation service based on the Alibaba Cloud's years of best security practices, which is provided for cloud users by Alibaba Cloud security engineers. It aims to build and continuously optimize the cloud security defense system to guarantee the business security.

Analysis

The process of converting strings to terms. Depending on the used analyzer, phrases FOO BAR, Foo-Bar, and foo,bar may be converted to the terms foo and bar. These terms are stored in the index. Perform a full text index query (not a term query) on FoO:bAR. This phrase is also converted to the terms foo and bar, and matches the terms stored in the index.

Analytical task

A GAS analytical algorithm task that is submitted by using GAS SDK. It is also called GAS UDF task.

API reference

Besides widely known as the Application Programming Interface, API also refers to the documentation of an API, or a help document that provides instructions of request parameters and request samples for developers to quickly understand and use it.

Apsara

A hyper-scale and universal computing operating system that is independently developed by Alibaba Cloud and provides services for the world. It can connect millions of servers all around the world to form a supercomputer and provide computing capabilities for the community in the form of online public services. From the PC-based Internet to the mobile Internet and Internet of Everything, the Internet has become a new infrastructure in the world. Apsara intends to resolve the scaling, efficiency, and security problems in human computation. Apsara revolutionizes the three aspects of cloud computing to provide powerful, universal, and beneficial computing capabilities to everyone.

Apsara Name Service and Distributed Lock Synchronization System

A distributed coordination service that performs name resolution and lock synchronization. It runs Paxos-like protocols and consists of multiple servers in a tree structure, similar to a file system, to store data. It provides the capabilities to handle high-availability and high-concurrency user requests. As an important basic module in the Apsara system, Apsara Name Service and Distributed Lock Synchronization System offers distributed consistency services. Compared to traditional distributed consistency service products, it has better performance and supports horizontal expansion, which is parallel to ZooKeeper in the industry.

ApsaraDB for MongoDB cluster version

The cluster version of ApsaraDB for MongoDB. You can purchase multiple mongos and shard nodes, and combine them with a single ConfigServer to form a cluster version of ApsaraDB for MongoDB. This makes it easy to create a MongoDB distributed database system.

ApsaraDB for Redis

An online storage service that adopts a key-value pair format and is compatible with the open-source Redis protocol. It supports multiple data types, including string, list, set, SortedSet, and Hash, and provides advanced functions such as transactions, message subscription, and message publishing. With the storage method “memory + hard disk”, it provides high-speed data read and write capabilities and meets the requirements of data persistence.

Attach data disk

A method used in the Linux environment to attach a data disk first before reading the disk data. This method is not required in the Windows environment because you can directly read the disk data by connecting to the SATA interface or SAS interface.

Attribute

Data of the row where properties reside. No limit for the number of property columns in each row.

Aurora

The HA (high-availability) management system of Relational Database Service (RDS). It checks the status of the master database every three seconds. When it detects that the master database is down, it switches the SQL requests of users to the slave database quickly.

Authoritative DNS resolution

A service that resolves root domain names, top-level domain names, and domain names of other levels.

Authorization

A process in which permissions required to perform a certain job are granted to the related personnel. It is the key to organization operations.

Authorization policy

A set of language specifications defined by Alibaba Cloud to describe permissions.

Authorized object

An attribute of a permission group rule, representing the target to be applied. In Virtual Private Cloud (VPC), it can be a separate IP address or Classless Inter-Domain Routing (CIDR) block. In a classic network, it can only be a separate IP address (generally the intranet IP address of an Elastic Compute Service (ECS) instance).

Automatic snapshot policy

A rule that defines when the automatic snapshots are taken and how long they are retained.

B

Backend server

A set of cloud servers that accept load-balanced distribution requests. Server Load Balancer forwards access requests to this set of servers based on the rules you set.

Back-to-origin error

An error that occurs when the origin site fails to respond to the request forwarded by Alibaba Cloud Anti-DDoS Pro because the IP address of the origin site is blacklisted by the firewall.

Basic cloud disk

Disks installed on the cloud server to support low I/O loads. It provides Elastic Compute Service (ECS) instances with the I/O performance of hundreds of IOPS.

Bastion host

A host implemented based on the protocol forward proxy. It records data streams of common Operation & Maintenance (O&M) protocols, such as SSH, Windows remote desktop, and SFTP. Then, it replays the records by reassembling the protocol data streams for O&M auditing.

Behavior analysis

A method that displays time-based distribution frequency of an event.

Block storage

A low-latency, persistent, high-reliability, and random block-level data storage provided for Elastic Compute Service (ECS) by Alibaba Cloud. It supports automatically copying your data in a zone, which avoids data unavailability caused by unexpected hardware faults and protects your business against the risk of component faults. Similar to a disk, you can format, create a file system, and persistently store the data on the block storage mounted to an ECS instance.

Brute force password cracking

A type of attack used by attackers who keep trying to crack the password by continuously specifying the password ranges based on the password naming rules until all verifications are completed.

Buffer pool

A pool of database connections. It allows applications to reuse database connections that already exist in the pool, and reduces repeated attempts to create new database connections. This technology improves the system performance and avoids the costs of building new connections.

C

Callback

A method used to process the log sending results, such as successful sending and exceptions. You are not required to inherit ILogCallback if you select not to process the result.

Capacity planning

A complicated and continuous performance study on resource consumption of softwares and hardwares in an existing system. Performance stress testing becomes a daily routine. Users in charge of applications can easily view application performance metrics and perform precise capacity planning for applications based on those metrics and the current system operating water level.

Channel muted

An alarm is not retrigged within 24 hours if the metric continuously exceeds the threshold value.

Classified protection

A basic information security protection system in China. It motivates citizens, legal persons, and other organizations to grade information systems for security protection based on the nationally formulated and consistent administrative norms and technical standards for classified protection of information security. Security levels (1-5) are assigned (low to high) based on the importance of the information system. Different security levels have different protection strategies. Cloud services feature reliable quality and security, and have passed the certification inside and outside of China, such as Trusted Cloud Initiative, Regulations of the People's Republic of China for Safety Protection of Computer Information Systems and Administrative Measures for the Graded Protection of Information Security, Information Security Management Systems (ISMS), and CSA STAR Certification.

Classless Inter-Domain Routing

A method used to assign IP addresses to users and efficiently route IP packets on the Internet to classify the IP addresses.

Cloud disk

An independent disk that can be mounted to any Elastic Compute Service (ECS) instance in the same region and zone. It can be classified into three types based on the performance: Ultra cloud disks, SSD cloud disks, and basic cloud disks.

CloudMonitor

A centralized monitoring platform for Alibaba operation engineers and developers, which monitors the devices, networks, databases, application clusters, and core services of Alibaba Group.

Column storage

An aggregate storage that holds data of the same fields in all records.

Command line

A command that runs in DOS.

Common neighbor analysis

An analysis on common contacts of a group of objects of the same type or different types.

Computing resource

A resource that can execute algorithms and calculations, such as E-MapReduce. A standalone system is also a computing resource. For example, YARN (Hadoop) and a standalone system with Tensorflow or Linux Shell are all computing resources.

Computing unit

A unit of computing resources.

Connection address

The host address that connects to the cloud database. It is displayed as a domain name and can be queried in the console.

Connection password

The password used to connect to ApsaraDB. For ApsaraDB for Redis, the password is in the format of instance ID:custom password. For example, if you set the password to 1234 when purchasing an ApsaraDB for Redis instance and the allocated instance ID is xxxx, the connection password is xxxx:1234. The connection password for other versions of ApsaraDB can be set in the console.

Consumer

Each consumer is assigned with several shards and must consume the shard data. The consumer name in the same consumer group is unique.

Consumer group

A sub-resource of Logstore. Consumers with the same consumer group name consume data of the same Logstore, but the data consumed by each consumer is different. Up to five consumer groups can be created under one Logstore, and the group names must be unique within the Logstore. Different consumer groups under the same Logstore consume data independently.

Consumer heartbeat

A request that consumers must regularly report to the server to indicate the normal status.

Content moderation

Based on the deep learning technology and Alibaba's years of big data analysis, it provides the intelligent risk identification service for pictures, videos, texts, and other multimedia contents. It not only reduces the pornographic, violent, terrible, and political contents, and other illegal or inappropriate contents that can cause violation risks, minimizes invasive advertising promotion, abuse, and other risks that affect user experience, but also greatly reduces manual auditing costs.

Co-occurrence analysis

An analysis method that searches multiple time-space domains for the objects that have footprints in more than one time-space domain. It is used by law enforcement to screen suspects of joint cases. It supports the following co-occurrence analysis models: circular, rectangular, and polygonal.

Cookie

Data (usually encrypted) that websites store on your local terminal for identity identification and session tracking.

Creating

The running status of a software or system. It indicates the software or system is being created.

Custom image

Created by using snapshots. The created image can be used to create an Elastic Compute Service (ECS) instance.

D

Dashboard

- A monitoring that provides quick and intuitive data models and other detailed metrics.
- A function in CloudMonitor that allows you to view monitoring data in a custom way. It is a centralized display of product instances for the same business and allows you to view monitoring data across products and instances.

Data access layer (DAL)

The persistence layer that is responsible for database access, such as reading from the database and transmitting data.

Data application layer

An abstraction layer that defines and generates personalized and diversified data metrics applicable to multiple business requirement scenarios.

Data disk

A disk that contains only data, excluding the operating system.

Data import

The process where data is migrated from one system to another system in batches. Migrated data includes structures, data, and definitions.

Data intermediate level

A bridge between the base layer and application layer, also known as the common data center. It abstracts common topics from all business data to develop standardized and shareable data.

Data mid-end

A framework and model of data construction, management, and application. It shields the complexity of business systems in the underlying hardware, processes and outputs standard, diversified, and converged data, and provides a variety of convenient data applications to support different business scenarios in the upper layer and drive business intelligence by using intelligent data.

Data owner

The person in charge of a database or a table. They are responsible for the management and control of other users who use the database or table.

Data persistence

A method to convert data between data models in the memory and storage models.

Data protection mechanism

Data protection can be categorized by technology. The technologies involve two aspects: the mechanism and the components or entities to be executed. The mechanism describes how to implement protection and where to execute components. The mechanism is divided into external data protection, data transformation, and internal data protection based on the technologies it depends on.

Database

- A database in ApsaraDB for Redis. ApsaraDB for Redis supports 256 databases. Data is written to database 0 by default.
- A warehouse that organizes, stores, and manages data based on the data structure. It includes tables, indexes, and metadata of database objects.

Database engine

The core service used to store, process, and protect data. It provides controlled access and rapid transaction processing to meet the requirements of enterprise applications that process massive data. It can be used to create relational databases that process transactions and data analysis online. This includes the creation of tables that store data and database objects that are used to view, manage, and guarantee data security (such as indexes, views, and storage procedures).

Database throughput unit (DTU)

A measuring unit that combines CPU, memory, and read/write rate, representing the performance of database engines.

Debugger

A computer program or tool that is used to debug other programs. The code to be examined can be alternatively run on an instruction set simulator (ISS) to check the running status and troubleshoot errors. This technology is useful when development is bottlenecked or problems cannot be located.

Decoupling

A process that converts mathematical equations with multiple variables to an equation group with a single variable. This simplifies analytic calculation as the result of the equation is not directly

affected by multiple variables. A multi-variable system can be converted to multiple independent single-variable systems by controlled variable selection and coordinate transformation, which means decoupling variables.

Deep learning

A branch of machine learning. It is an algorithm that attempts to use multiple processing layers with complex structures or composed of multiple nonlinear transformations for high-level abstraction of data.

Denial of Service (DoS) attack

A cyber-attack that attackers send massive data packets to the target server to consume the network bandwidth and computer resources of the target computer, which prevents the target computer from providing services to users normally.

Deployment set

Capabilities provided by Elastic Compute Service (ECS) that allow you to detect the physical topology of the host, rack, and switch and select a deployment policy that meets the business requirements according to your business type to improve the overall reliability or performance of your business.

Destination Network Address Translation

The action of translating the destination address of an IP packet to another address.

Detailed list

A list that contains attributes of relations involved in one or more objects.

Developer

The resource consumer and data constructor in a project, who can edit and delete resource objects from a project.

Development and testing/production environment

Two environments that every workspace has. They are two separate resource configurations. Algorithm engineers can register algorithms, configure algorithm templates and processes, and test them only in the development and testing environment. After the test is passed, the algorithm processes in the testing environment must be released to the production environment by "bringing them online". No one can edit the algorithm processes in the production environment, but the parameters can be modified and all the modifications are recorded in logs.

Distributed

A mode in which different services are distributed on different areas and devices, and each node performs a specific task.

Distributed Denial of Service (DDoS) attack

A cyber-attack that an attacker sends a large number of valid requests to make the network become unavailable by consuming huge network resources. The attack can be divided into the following types: Attacks that interrupt or even block normal network communication by overloading the network. Attacks that overload servers by submitting massive requests to the servers. Attacks that prevent an authorized user from accessing a server. Attacks that prevent a service from communicating with a specific system or individual.

Distributed System Performance Monitor

The module responsible for information collection, monitoring, and diagnostics. It deploys a lightweight information collection module on each physical machine to obtain the running status of the operating system and applications. It monitors cluster faults and analyzes engines to evaluate the running status of Apsara system.

Docker image

A standard packaging format of a container application. You can specify an image to deploy containerized applications. The image can be from Docker Hub, Alibaba Cloud Container Hub, or your private registry. An image ID is uniquely identified by the URI of the image repository and the image tag (latest by default).

Domain name

The name of a computer or a computer group on the Internet. It consists of several parts separated by periods (.) to identify the electronic location of the computer or computer group for data transfer.

Domain name resolution

A function of mapping a domain name to a corresponding IP address based on DNS.

Downstream traffic

The number of bytes that are downloaded from the Internet.

Dry run of scheduling

An operation that starts a task at a specified time point during data scheduling, but does not run task scripts. The task is considered successfully run by default.

E

ECS Bare Metal Instance

A new computing server product that combines the elasticity of virtual machines and the performance and features of physical machines, which is built based on the next-generation virtualization technology that is independently developed by Alibaba Cloud.

Elastic Compute Service (ECS)

A simple and efficient cloud computing service that features elastic processing capabilities. It supports many operating systems, such as Linux and Windows.

Elastic Network Interface (ENI)

Also known as secondary ENI. You can create a secondary ENI and attach/detach it to/from an instance. The ENI created by default when the Virtual Private Cloud (VPC) instance is created is a primary ENI. The lifecycle of the primary ENI is the same as that of the instance. You cannot detach the primary ENI from the instance.

Endpoint

The access domain name that Object Storage Service (OSS) provides external services. OSS provides external services in the form of HTTP RESTful API. Different domain names are required when you access different regions. The endpoint varies if you access the same region by using Internet or intranet.

Envelope encryption

The action of generating a symmetric key for data to be encrypted and encrypting the symmetric key by using a specific master key, which keeps the symmetric key protected in an envelope. The key is protected in the envelope in unsecured communication processes, such as data transmission and storage. The key is retrieved from the envelope only when the symmetric key must be used.

Ephemeral disk

A disk (hard disk) that is installed on the device board where other disks are located. Do not insert, remove, or move the disk without permission. Generally, the disk includes the partition where the computer operating system resides and other partitions.

Ephemeral SSD

A Solid-State Disk (SSD) that is installed on the device board where other disks are located.

Unlike the basic cloud disk, data is not triplicated but the Redundant Array of Independent Disks (RAID) technology is used. Therefore, you cannot avoid the data loss if a single physical machine is down.

Eviction policy

The same as the eviction policy of ApsaraDB for Redis. For more information, see <http://redis.io/topics/lru-cache?>.

Expansion

A method used to scale up the capacity of an Elastic Compute Service (ECS) instance disk.

Express Connect

A data transmission channel based on the high-quality infrastructure provided by Alibaba Cloud. It is used in different private networks for secure and reliable intranet communications, such as between two Virtual Private Cloud (VPC) instances and between a VPC instance and an on-premises data center.

F

Field Programmable Gate Array (FPGA) instance

A computing instance that provides the FPGA, which is based on the Alibaba Cloud elastic computing framework and allows you to create FPGA instances and custom dedicated hardware accelerators in minutes.

Flow cleaning

A service used to monitor the data flow coming into your Internet Data Center (IDC) in real time and identify the abnormal flow, such as DoS attacks, in time. Abnormal flow is cleaned without affecting the normal services.

Forward compatibility

A design characteristic that allows earlier versions to support data generated by later versions.

Forwarded domain name

A domain name that is resolved by the local DNS server. Domain name resolved data is configured and managed on the local DNS server. The local DNS server forwards domain name resolution results to another DNS server for further resolution.

Forwarding weight

The weight for traffic forwarding that you can specify for each Elastic Compute Service (ECS) instance in the backend server pool according to the external service capabilities and situations of the backend ECS instance. The higher the weight is, the more access requests are assigned to the ECS instance.

FreeBSD image

An Elastic Compute Service (ECS) instance environment based on the FreeBSD operating system. Currently, the supported versions include 10.1 (64 bit).

Full backup

The process of backing up all data or applications at a time point.

Full snapshot

The first created snapshot of a disk that saves all the disk data.

G

Geo bubble chart

A rendering style where circles with defined colors and sizes are positioned on specified locations on a geo chart.

Geo chart

A map of a country, a continent, or a region with areas identified based on specified data.

GetHistograms

A method used to query the number of logs that meet the query condition, and the distribution of these logs in the queried time range (namely, the number of logs that meet the query condition in each sub time range), without returning the specific raw logs.

GetLogs

A method used to query logs that meet the query condition, with all raw logs returned.

Globally unique ID

A numerical ID that is globally unique. It helps keep data such as unique keys and primary keys globally unique (on all nodes) in a distributed environment.

GPU instance

A GPU-based computing service featuring timeliness, high speed, and strong capabilities of parallel computing and floating-point computing. It mainly applies to the scenarios such as video decoding, graphics rendering, deep learning, and scientific computing.

Graph instance

A set of computing, storage, and metadata nodes that provide graphing services. It is a collection of resources and data and also the minimum service unit of graph computing.

Graph instance management

Management of graph instances. You can manage graph instances on the user terminal or the operation terminal. Generally, you can authorize users and view authorization code on the user terminal, and create, stop, delete, and update the graph instances online on the operation terminal.

Group analysis

A method that analyzes the relationships among objects of the same type or different types. The relationships include direct and indirect relationships.

Guest

A participant of a project who can view resource objects of a project, but cannot edit or delete them.

H

Health check

A basic function of Server Load Balancer, which is used to check the business availability of the backend servers (Elastic Compute Service (ECS) instances), improve the overall availability of the frontend business, and avoid the effect on the overall services caused by the abnormal backend ECS instances.

High concurrency

A factor that must be considered in the distributed system architecture design of Internet. It uses a method to make sure that the system can handle large numbers of requests concurrently.

High-availability cluster

A server cluster technology that reduces service interruptions. It protects your business programs to provide uninterrupted services and minimizes the impact caused by softwares, hardwares, or human errors.

High-Availability Virtual IP (HAVIP)

An intranet IP resource that can be created and released independently. You can use the ARP protocol on the Elastic Compute Service (ECS) instances to declare the IP.

Host intrusion detection

A method that Alibaba Cloud Security uses to perform security analysis on the host logs, detect and report the aggressive behaviors in real time, such as system password cracking and logons from abnormal IP addresses.

HyperText Transfer Protocol

A universal, stateless, and application-oriented protocol for distributed, collaborative, and hypermedia information systems.

I

ID

The ID used to identify a document. The index, type, and ID of a document must be unique. The ID is automatically generated if not provided.

Image processing service

An extended service of Object Storage Service (OSS). It is built on OSS and provides real-time image processing functions, such as resizing, cropping, reformatting, and adding watermarks to images. You must enable image processing in the OSS console first, and then call the RESTful interface provided by the image processing service to process the images stored on OSS, and obtain the processed images immediately.

Image sharing

The action that allows you to share your custom image with other Alibaba Cloud accounts. Then, these accounts can use the shared custom image to create instances and change the system disk.

Index

A table in a relational database. It contains a mapping that defines different types of index fields. An index is a logical namespace that maps to one or more primary shards and zero or more replica shards.

Information cube

Behavior analysis: Displays time-based distribution frequency of an event.

Behavior details: Displays details about an event (original data records are filtered by rule).

Object information: Displays information about entities in relational networks and groups the entities by type.

Statistics: Displays statistics about relationships and entities in relational networks, including overall distribution, object familiarity, and relationship attributes.

Infrastructure as Code

The process where you can perform operations on Alibaba Cloud resources by using templates, APIs, and SDKs, which is similar to writing codes.

Ingress

A collection of rules that authorize the inbound access to the cluster. You can provide the externally accessible URL, Server Load Balancer, SSL, and name-based virtual host by using the Ingress configurations. You can request Ingress by posting Ingress resources to API servers. Normally, the Ingress Controller is used to implement Ingress by using Server Load Balancer. It can configure the edge router and other frontends, which helps you handle the traffic in the HA method.

InnoDB

A binary standard database engine of MySQL. Based on traditional ISAM and MyISAM engines, InnoDB also supports the ACID-compatible transaction function.

Instance

An independent resource entity that contains basic resource elements.

Intelligent data R&D

The system can perform optimal data calculation and storage based on user input to automatically implement data R&D.

Internet address

A connection address used by a source IP address to provide services for the Internet client.

Internet Gateway

A gateway for access to the Internet.

Intranet

The network within an organization. It is usually a local area network (LAN). A LAN is a computer communication network that interconnects computers, external electronic devices, and databases within a limited geographical area. The LAN is connected to the databases, LANs, or processing centers in other regions by using data communication networks or dedicated data circuits, which forms an information processing system that covers a wide geographic area.

J

Job Scheduler

A distributed resource management and job scheduling system. Parallel to Kubernetes in the industry.

Job scheduling

An important part of the operating system. For real-time operating systems, it directly affects their real-time performance.

K

Keepalived

An open-source software that guarantees the high availability of clusters and detects the status of Web servers. If it detects that a Web server is down or has an error, the Web server is automatically removed from the system. The removed Web server is automatically added after resuming the normal work.

Kubernetes

An open-source and large-scale container orchestration and scheduling system of Google. It is used to automatically deploy, expand, and manage containerized applications and has the characteristics such as portability, scalability, and automatic scheduling.

Kubernetes service

The basic operation unit of Kubernetes. An abstract of real application services, each service has multiple containers to provide the support. The port of Kube-Proxy and service selector determine if the service request is forwarded to the backend container, and a single access interface is displayed externally. You are not required to know how the backend works, which brings great benefits to backend expansion and maintenance.

Kubernetes volume

Similar to Docker volumes. The only difference is that the scope of Docker volumes is a container, while the lifecycle and scope of Kubernetes volumes are a pod. The volumes declared in each pod are shared by all the containers in the pod. You can use Persistent Volume Claim (PVC) logical storage to ignore the actual backend storage technology. The specific configurations about Persistent Volume (PV) are completed by the storage administrator.

L

Layer-4 Server Load Balancer

A service that uses the information defined at the network transport layer (layer 4) as the basis for deciding how to distribute client requests across a group of servers. For Internet traffic, it makes the load balancing policy based on the source and destination IP addresses and ports recorded in the packet header, without considering the contents of the packet.

Layer-7 Server Load Balancer

A service that provides load balancing for services at the application level, such as HTTP services. It determines the routing based on the characteristics of the HTTP header and the actual contents of the message, such as the URL, data type (text, video, and graphics), and information in the cookie.

Lineage analysis

A method that presents the lineage relationships among all family members.

Linux virtual server

A cluster technology that supports the IP load balancing and content-based request distribution techniques.

Log

An abstraction of system changes during the running process. The minimum data unit processed in Log Service. Its content is a time-ordered collection of some operations and the corresponding operation results of specified objects. LogFile, Event, BinLog, and Metric data are its different carriers. In LogFile, each log file is composed of one or more logs, and each log describes a single system event.

Log group

A collection of logs. The basic unit for reading and writing logs. Its maximum capacity is 4096 logs or 10 MB.

Log Service

A one-stop service for log data, which allows you to quickly complete the collection, consumption, shipping, query, and analysis of log data without the need for development, improve the Operation & Maintenance (O&M) efficiency and the operational efficiency, and build the processing capabilities to handle massive logs in the data technology (DT) era.

Log topic

A topic used to classify logs in a Logstore, which can be specified when you write and query logs.

LogHub

The basic function of Log Service, which provides the real-time collection and consumption of logs.

LogHub Client Library

An advanced mode provided for the LogHub consumers, which is used to solve the issue of automatic shard allocation when multiple consumers consume Logstores simultaneously. For example, in this mode, logics such as shard load balancing and consumer failover are automatically processed among multiple consumers in the Storm and Spark scenarios. This mode contains four concepts: consumer group, consumer, heartbeat, and checkpoint.

LogHub Log4j Appender

Log4j is an open-source project of Apache, which allows you to set the log output destination to console, file, GUI component, socket server, NT event recorder, or UNIX Syslog daemon. You can also set the output format and level of each log to control log generation with a finer granularity. These settings can be configured flexibly by using a configuration file, without modifying application codes. Log4j is consisted of three important components: log level, log output destination, and log output format. The log level includes ERROR, WARN, INFO, and DEBUG in descending order of priority to indicate the importance of logs. The log output destination indicates whether logs are printed to the console or a file. The log output format indicates the displayed content of logs. LogHub Log4j Appender allows you to set the log output destination to Alibaba Cloud Log Service.

LogHub Producer Library

A LogHub class library written for high-concurrency Java applications. Producer Library and Consumer Library are the read and write packaging for LogHub to lower the threshold for data collection and consumption. LogHub Producer Library solves the following issues: Logs collected from the client are not flushed into the disk. Data is directly sent to Log Service by using the network after being generated. High concurrency write operations on the client. For example, more than one hundred write operations are performed in one second. Client computing is logically separated from I/O. Printing logs does not affect the computing time used. In the preceding scenarios, Producer Library simplifies your program development steps, aggregates write requests in batches, and sends the requests to the LogHub server asynchronously. During the process, you can configure the parameters for aggregation in batches and the logic to process server exception.

LogSearch

The action of real-time index and data query, which allows you to query logs based on time and keywords after creating indexes for LogHub data.

LogShipper

The action of shipping LogHub data to storage services for storage and big data analysis.

Logstore

A unit of Log Service for collecting, storing, and querying logs. Each Logstore belongs to a project and each project can create multiple Logstores.

Logtail

A log collection client provided by Log Service that facilitates the log access. By installing it on your machine, you can monitor the specified log file and have the logs newly written to the log file automatically uploaded to your specified Logstore.

Logtail configuration

A configuration that defines how to collect a type of logs from the machine, and then parses and sends the collected logs to the specified Logstore. You can add the configuration for each Logstore in the console and then the Logstore can receive the logs collected with the configuration.

M

Marquee analysis

The search for objects based on user-defined geographical locations and ranges of space and time.

Massively parallel processing

A distributed shared-nothing computing architecture, where each node is completely independent of other nodes and multiple nodes perform concurrent computing to improve performance.

Match tracks

An analysis that calculates unknown objects that have similar trajectories with known objects within specific time periods. It is used to identify crime associates.

Measurement

A numeric value related to measurement, comparison, or aggregation. For example, a sales revenue field is a measurement because you can average and total the data.

Memory

Internal storage where the computing data in CPUs and the data exchanged with external storages, such as devices and hard disks, are temporarily stored.

Merge

The process in which the incremental data in the memory is combined with the baseline data on the disks. The combined data is written to the disks as new static data, and the memory is released.

Message

A data unit exchanged and transmitted in a network, that is, a data block to be sent by a site at one time.

Message key

The business identifier of a message, which is set by the message producer and uniquely identifies a business logic.

Metadata

Also known as data about data. It defines data properties.

Metadatabase

The database where data is organized, stored, and managed based on the data structure.

MetaService

A service that allows you to access Alibaba Cloud resources in the E-MapReduce cluster without using AccessKey.

Monitored item

Monitored data types set by users or the system by default. For example, the HTTP monitoring of site monitoring has two monitored items by default: response time and status code. For Elastic Compute Service (ECS), the monitored items include CPU utilization and memory utilization.

Monitoring system

A system that checks if the system is functioning properly, collects status and performance data of instances, such as monitoring the instances, physical resources, resource usage, and network, and sends alarm notifications.

Multipart upload

A method of uploading a large file by separating it into smaller parts.

Multi-Service Transmission Platform

A unified network management platform for multi-service nodes based on SDH, which supports access, processing, and transmission of TDM, ATM, Ethernet, and other services.

Multi-zone

An area composed of multiple zones, which is deployed in different regions by Alibaba Cloud Server Load Balancer to achieve disaster recovery across data centers in the same region. If the

data center in the active zone malfunctions or becomes unavailable, Server Load Balancer rapidly switches to the data center in the standby zone to restore its service capabilities within 30s. When the active zone is restored, Server Load Balancer automatically switches back to the data center in the active zone to provide services.

N

Namespace

A set of unique names that provide virtual isolation for the Kubernetes cluster. The Kubernetes cluster starts with 3 initial namespaces: the default namespace and two system namespaces (kube-system and kube-public). Administrators can create namespaces as required.

NAT Gateway

A gateway for cloud service access.

Network device model library

A library that contains the network device specifications, including model name, hardware configuration, and purchasable vendor model, which are verified by the Apsara Stack console and can be officially output.

Node

A server (either an Elastic Compute Service (ECS) instance or a physical server) that has a Docker Engine installed. It is used to deploy and manage containers. The Agent program of Container Service is installed on a node and registered to a cluster. The number of nodes in a cluster is scalable.

Normal scheduling

An operation that runs task scripts during data scheduling.

Number of processes

The number of processes. A process is a computer program activity that is run for a data set. It is the basic unit for resource allocation and scheduling, and also the foundation of the operating system structure.

O

Object Storage Service (OSS)

A massive, secure, cost-effective, and high-reliability cloud storage service provided by Alibaba Cloud. You can call APIs to upload and download data from any application anytime and anywhere, and perform simple data management in the Web console. OSS can store all types of files and can be used by websites, development enterprises, and developers.

Offline computation

A computing process that is performed after the generated data is stored.

Open source

A certification mark registered by and an official definition given by the Open Source Initiative (OSI). Softwares that bear this mark are open source. An open source software can be used, modified, and released without license restrictions, and its source code is available for public use.

Operation log

A record of all operations, including the time and content of each operation.

Orchestration service

A service that manages a group of cloud resources in a centralized way by using logical sets such as resource stacks. Therefore, cloud resources can be created, deleted, and cloned in resource groups. In DevOps practices, the development, testing, and online environments can be easily cloned. It simplifies the overall migration and expansion of applications.

Orchestration template

A text file that you can read and edit with ease. You can control the infrastructure version by controlling the orchestration template version with tools such as SVN and Git. You can achieve Infrastructure as Code (IaC) by integrating the orchestration capability of ROS with your own applications with APIs and SDKs. The orchestration template also provides a standard way of delivering resources and applications. You can use the template to deliver integrated systems and solutions that include cloud resources and applications.

P

Partition key

The first primary key column in a primary key. Rows with the same partition keys belong to the same data partitions. A single data partition can contain multiple partition keys.

Partitionally ordered message

The message partitioned by a shard key. For a specified topic, all messages are partitioned based on the shard key. Messages in the same shard are released and consumed in strict FIFO order.

The shard key is a key field used in ordered message to distinguish different shards, which is completely different from the key used in normal messages.

Password-free access

You can access ApsaraDB for Memcache on an authorized Elastic Compute Service (ECS) instance without providing your username or password. For more information, see Password-free access in ApsaraDB for Memcache User Guide.

Peak bandwidth

The maximum permitted traffic burst that a site can sustain.

Permission group

The whitelist provided by NAS. You can add rules to a permission group to allow IP addresses or Classless Inter-Domain Routing (CIDR) blocks to access the file system with different permissions. Each mount point must be bound to a permission group.

Persistent connection

A connection used to send multiple data packets continuously. During the connection, both parties must send link detection packets if no data packets are sent.

Persistent volume claim (PVC)

A consumer of storage resources. It varies with business service requirements. It is configured by the user of the Kubernetes cluster, who is the service administrator.

Physical backup

Dump of physical files, such as data files, control files, and archived logs. These files can be used to restore abnormal databases.

Plug-in

The system built-in plug-ins used to process universal data during the import process. You can select a plug-in by using content conversion when defining application structures or configuring data sources.

Port forwarding

A method provided by SSH for secure network communications. It forwards the behavior of a network port from one network node to another. It allows an external user to access a port of a private internal IP address from the external by passing an activated NAT router.

Pre-check

An inevitable stage before the migration task is started. It checks the preconditions that will affect migration, such as the connectivity between the source instance and target instance, and permissions of the account to be migrated. If the pre-check fails, make the appropriate changes and try again.

Primary key

The unique identifier of each row in a table. It consists of one to four columns. When a table is created, the primary key must be defined. To define a primary key, make sure that the name and data type of each primary key column, and the fixed sequence of the primary key columns are available. The data types of primary key columns can only be String, Integer, and Binary. If the data type is String or Binary, the size of the column value cannot exceed 1 KB.

Primary shard

Each document is stored on a single primary shard. When you index a document, it is indexed on the primary shard first and then on all replicas of the primary shard. An index consists of five primary shards by default. You can add or remove primary shards based on the number of documents that the index can handle. After an index is created, the number of primary shards cannot be changed.

Primary-secondary dual node instance

An ApsaraDB instance that adopts a primary-secondary architecture. The capacity and performance that primary-secondary dual node instances can expand are limited.

Private key

The secret key in a key pair that is generated by using an algorithm. A key pair contains a private key and a public key. The public key is the public part and the private key is the private part of a key pair. Generally, the public key is used to encrypt session keys, verify digital signatures, or encrypt data that can be decrypted by the corresponding private key. The key pair generated by using the algorithm is globally unique. Data encrypted with one key in the key pair can only be decrypted with the other key of the same key pair. For example, data encrypted with the public key must be decrypted with the corresponding private key, and data encrypted with the private key must be decrypted with the corresponding public key.

private-read-write

One of the three access permissions to Object Storage Service (OSS) buckets and objects. The other two access permissions are public-read-write and public-read. The private-read-write permission only allows the bucket creator to read and write files in the bucket, while other users cannot access the files in this bucket.

Product baseline

Metadata about how cloud products are deployed on Apsara Stack. It is a set of standard models that contain software components and hardware compatibility. The Apsara Stack business team sets the standard formats, in which product designers and developers of each product line provide the contents. After these formatted contents are accepted and versioned by Shuguang team, they become official product baselines. Each version of Apsara Stack has its own set of product baselines.

Product summary

The overview and summary about what a product is, what it does, and how it works.

Project

- After entering a project, you can directly perform operations on all objects in it.
- The most basic business organizational unit of Alibaba Cloud StreamCompute, and the basic unit used for managing clusters, jobs, resources, and staff.

Project administrator

An administrator of the data project.

Project version

The planning of a project usually involves the comparison of different schemes, which then evolve into a unique final scheme. Therefore, each project version corresponds to a suite of independent planning and implementation scheme. All the versions available are called optional versions, which are presented for selection in the pre-sales project management system. The version finally selected is called the locked version.

Property

Characteristics in an object relation. In I+, a property belongs to a relation or entity and cannot exist independently. For example, properties of a person include height, weight, birthplace, and name. Properties of a mobile phone number include the home location and operator. A primary key is also a property. For example, an ID number is a property of a person, and a mobile phone number is a property of a mobile phone.

Public data center

The data center that abstracts public topics from all service data to develop standardized and shareable data.

Push

A proactive action to send resources from the origin site to edge nodes. You can directly hit cache when accessing resources, which lightens the pressure of the origin site caused by bursts in the traffic back to the origin.

R

Rack

A physical framework that is designed to place devices required in a data center, such as servers, network devices, and cables. This construction is usually used in a data center.

RDS Operations and Maintenance System

The O&M system of Relational Database Service (RDS).

Read and write operations per second

The number of times that an application reads from and writes to a database per second.

Read-only SQL

A type of SQL statement for reading data, such as SELECT statements.

Redis

A high-performance Key-Value storage system (cache and store) released in compliance with the BSD open-source protocol.

Reference manual

A help document for users. It contains all reference contents that users need to know when they use the product.

Relational database

A stable, reliable, and elastically scalable online database service. Based on Apsara distributed system and high-performance storage, Relational Database Service (RDS) supports the MySQL, SQL Server, and PostgreSQL engines. It provides a complete set of solutions for disaster recovery, backup, recovery, monitoring, and migration to free you from worries about database O&M.

Relationship analysis

An analysis method that focuses on mining small numbers of relevant clues and intelligence from large amounts of unrelated information. It is a process of converting information into actionable intelligence. It helps you implement the unlimited association of information.

Release cluster

An operation that releases a cluster. You can release a cluster after offline jobs are completed. You can also dynamically add nodes.

Remote logon reminder

The alarm triggered if an unfamiliar remote logon is detected by Server Guard. Server Guard regularly collects the logon logs of hosts to the cloud, analyzes and compares the logs on the cloud, and triggers an alarm if the remote logon is detected.

Replica

A replica that stores data for a partition. Multiple replicas are provided for each partition to guarantee the data security and provide high-availability data services. Each replica stores the following types of data:

- SSTable (static data that is stored in a disk)
- MemTable (incremental data that is stored in a memory)
- Logs that record transactions

Replicas are automatically distributed to multiple servers based on load balancing and specific policies. A replica can be migrated, copied, added, deleted, or converted to a different format.

Reserved read/write throughput

An attribute of table. Table Store reserves certain resources for access. The access within the reserved read/write throughput consumes less resources.

Response time

The time from when a query or update statement sends a request to when data is received.

Role

A set of permissions. It is generated when a dimension is associated with foreign keys to produce a star or snowflake model. One role is generated every time a same dimension is associated with different tables. When a dimension is associated with two foreign keys in a logical table, different aliases are required to uniquely indicate the roles of this dimension in the table.

Round robin

A scheduling algorithm that Server Load Balancer uses to distribute the traffic. Requests are sequentially distributed to the backend servers.

Route computing

A computing method that reduces the routing overhead to improve routing protocol functions.

Route entry

Each item in the route table, which defines the next hop address for the network traffic destined to a specific destination Classless Inter-Domain Routing (CIDR) block.

Route table

A list of route entries in a VRouter.

Routing

When you index a document, it is stored on a single primary shard and mapped to a shard by hashing the routing value. By default, the routing value is derived from the document ID or ID of the parent document (make sure that the child document and parent document are stored on the same shard). This value can be modified by specifying a routing parameter during the index, or using a routing field in the mapping.

Row storage

An aggregate storage that holds the data of all fields of each record.

Rule engine

An online service for scenarios where the business rules change frequently. It helps you separate business rules from the application code. It combines predefined conditional factors to compile business rules flexibly and then makes business decisions based on the business rules.

Running

The normal operating status of instances. You can run your services on instances that are in this status.

S

Scheduling period

The interval for scheduling data.

Schema migration

A type of task migration. In database migration, it refers to the migration of syntax defined by schema objects, such as tables, views, triggers, stored procedures, storage functions, and synonyms. For migration between heterogeneous databases, data type mapping is carried out in the schema migration phase. The syntax defined by objects is adjusted based on the syntax definitions of source and target instances. For example, when Oracle is migrated to MySQL, the number in Oracle is mapped to the decimal type in MySQL.

Security group

A virtual firewall with the status detection and package filter functions. It is used to set network access control for one or more Elastic Compute Services (ECS) instances. Instances in the same security group have their networks interconnected, while instances in different security groups cannot communicate with each other by using the intranet by default. You can authorize the intercommunication between two security groups.

Security group inbound rules

A set of rules that enable or disable the Elastic Compute Service (ECS) instances associated with the security group to access the Internet or intranet from the inbound direction.

Security group outbound rules

A set of rules that enable or disable the Elastic Compute Service (ECS) instances associated with the security group to access the Internet or intranet from the outbound direction.

Separate mount

The action of mounting disks to any Elastic Compute Service (ECS) instance in the same zone.

Server

An application that agrees to request-side connections and sends responses. Applications can act as both a client and a server.

Server Guard

A sub-product of Apsara Stack Security, which acts as a server security O&M manager. With the linkage between the lightweight Agent plug-in installed on the server and rules of the cloud protection center, it detects and protects against the intrusion events in real time to guarantee the server security.

Server Load Balancer

A service used to distribute traffic for multiple Elastic Compute Service (ECS) instances. It can expand the external service capabilities of applications by distributing the traffic and improve the application availability by eliminating single points of failure.

Server Load Balancer instance

A running instance of Server Load Balancer, which must be created if you want to use Server Load Balancer.

Server Load Balancer listener

A concept of Server Load Balancer instances, which includes frontend ports, backend ports, Server Load Balancer policies, and health check configurations. Each listener corresponds to a backend application service.

Server Load Balancer listening

A function of listening to the ports, load balancing policies, and health check configurations.

Service

A set of containers based on the same image and configurations, which can be considered as a scalable microservice.

Session persistence

A basic function of Server Load Balancer that distributes access requests from the same client to the same backend server for processing.

Shard merge

The process of merging two shards connected to two adjacent key ranges into a single shard.

Shared block storage

A block-level data storage device that supports concurrent read/write access of multiple Elastic Compute Service (ECS) instances and features high concurrency, high performance, and high reliability. Its data reliability reaches 99.9999999%. A single shared block storage can be attached to at most four ECS instances simultaneously.

Short-lived connection

A connection created when both parties have data exchanges. The connection is closed after the data is sent. The data transmission occurs only once for each connection.

Single-connection

The mode in which only one client is connected to a database. In this mode, data of 1 KB, 10 KB, 100 KB and 800 KB is used to perform read/write operations on Memcached and OCS, and the average response time of Memcached and OCS are compared.

Site monitoring

A monitoring function that stimulates the visits of real users and monitors items such as availability, connectivity, and DNS resolution.

Situation Awareness

A service used to collect 20 types of raw logs and cybersecurity threat intelligence reports from enterprises, and use machine learning to reconstruct past attacks and predict potential attacks.

Slow query

A SQL statement query that is not responded within the specified time.

Small table copy

A solution that copies small tables to sub-databases in JOIN connection by using data synchronization to improve JOIN efficiency. With this solution, you can implement multi-node peer-

to-peer synchronization of tables that do not change frequently, such as the configuration tables and constant tables. This speeds up the join queries between those tables and other split tables.

Snapshot

The data backup of a disk at a certain time point, which has two types: automatic snapshots and manual snapshots.

Snapshot quota

A quota limit for each disk. Each disk in Elastic Compute Service (ECS) can create at most 64 snapshots.

Software Defined Networking

A method that separates the control level of network devices from the data level to flexibly control the network traffic and provide a good platform for the innovation of core networks and applications.

Source field

By default, the JSON document that you index is stored in the `_source` field, and returned by all get and search requests. This allows you to access the original object directly from search results, rather than querying the ID and then obtaining the document.

Specifications

Data transmission provides link specifications with different performances for you, with the number of simultaneous records as the measuring standard.

Split

The action used to divide an object. Shard split is an operation that splits a shard into two shards connected by a shard key range.

Spreadsheet

A spreadsheet that organizes your data into rows and columns in a clear, instantly readable format.

SQL injection

The action of injecting the SQL commands to the query strings used for Web form submission, domain name input, or page request to trick the server and run the malicious SQL commands.

SSD cloud disk

A Solid-State Disk (SSD) that is installed on the Elastic Compute Service (ECS) instance, oriented to the I/O-intensive applications, and provides stable and highly random IOPS performance.

SSH key pair

A new authentication method provided by Alibaba Cloud that allows you to log on to the Elastic Compute Service (ECS) instances remotely.

Starting

The status of an instance when it is started or restarted in the console or by using an API until the instance runs normally. An exception occurs if an instance stays in this status for a long time.

Statement

A syntactic unit that consists of a word or a group of syntactically associated words to express propositions, doubts, commands, wishes, or exclamations.

Static/dynamic resource separation

Separation of static and dynamic resources from Web pages or mobile applications.

Stop scheduling

An operation that stops data scheduling. Tasks are not started at the specified time point.

Stopped

The status of an instance after it stops normally. An instance in this status cannot provide external services.

Stopping

The status of an instance after the stop operation is performed in the Apsara Stack console or by using an API until the instance actually stops. An exception occurs if an instance stays in this status for a long time.

Storage

An operation that saves or records information or data in a specific format in a computer or to an external medium, and guarantees that the information or data is accessible.

Storage engine

The MySQL storage engine uses different technologies to store data to files (or memory). Each technology uses different storage mechanisms and index techniques, and has different lock

granularities. These technologies provide a wide variety of functions and capabilities. You can select different technologies to obtain higher speed or more functions, which in turn improves the overall application functions.

Storage resource

The resource that can store data, including offline storage resources, online storage resources, and stream storage resources. TableStore and Relational Database Service (RDS) are online storage resources. A standalone system can also be considered as a storage resource that stores data as documents.

StreamCompute unit

The minimum operational capacity used to run a StreamCompute job. It is the capacity to handle event streams within a defined CPU, memory, and I/O configuration.

Stress testing

A method to test system stability. It is generally performed outside of the normal system operating ranges to identify system limitations and potential risks.

System disk

A disk that contains an operating system.

T

Table

The most basic database object. Each table consists of several rows, and each row has the same predefined columns. SQL statements are used to add, delete, query, and modify tables. A primary key, which is unique in the data set of the table, is typically composed of several columns of a table.

TableGroup

A group that tables can belong to. It is a logical concept, independent of physical data files. It only affects the scheduling method of table partitions.

Taobao Distributed Data Layer (TDDL)

The data engine that makes database and table shards transparent to applications, and enables data replication between heterogeneous databases. It is based on centralized JDBC data sources and provides active/standby, read/write splitting, and dynamic database configuration functions.

Task

A job consists of tasks and their dependencies. Batch computing supports jobs that can be described in the form of a directed acycline graph (DAG). The dependencies between tasks can be specified only when a job is submitted, and cannot be modified after submission.

Tesla

O&M system of Table Store (formerly OTS).

Transactions per second

The number of times that SQL statements are executed and the number of processed transactions per second.

Transmission Control Protocol

A connection-oriented and reliable transport layer communication protocol that is based on byte streams.

Trigger

A method that the SQL server provides for programmers and data analysts to guarantee the data integrity. It is a special storage process related to table events. It is initiated by events, not called by programs or started manually. For example, when a table is inserted, deleted, or updated, a trigger is initiated. The trigger is typically used to reinforce data integrity and business rules.

Trojan

Any malicious computer program which hides its true intent from users. Attackers obtain the administrator account of a website by many methods, such as SQL injection, website sensitive files scanning, and sever vulnerabilities. Then, they log on to the website backend and recover or upload the vulnerabilities by backing up the database to obtain a webshell. With the obtained webshell, they modify the webpage contents by embedding malicious code into the webpage. They can also use a weak password to obtain the FTP of the server or website and then directly modify the webpage contents. When accessing webpages embedded with malicious code, users will automatically be redirected to a specific address or download the Trojan viruses.

Type

The type of a document, such as emails, user files, or weet files. The search API filters documents by type. Each index includes multiple types, and each type contains a set of fields. One index contains multiple type field names, and same type fields must have same mappings.

U

Ultra cloud disk

A cloud disk that combines Solid-State Disk (SSD) and Hard Disk Drive (HDD) as the storage media.

Unity3D

An integrated game development tool compatible with multiple platforms. A fully integrated and professional game engine. Developed by Unity Technologies, this tool allows a player to easily create various interactive contents such as 3D video game, architectural visualization, and real-time 3D animation. Its logs can be collected conveniently by using the Web Tracking function of Log Service.

URL forwarding

The action of automatically redirecting you to a preset address when you access a domain name. It includes explicit and implicit URL forwarding.

V

Verification

A process that verifies a result.

Vertical data center

A data center that stores raw data from different business databases.

Virtual Border Router (VBR)

A product mapping of the leased line that you have applied to access a VSwitch. It can be regarded as a VRouter between Customer-Premises Equipment (CPE) and the Virtual Private Cloud (VPC), which serves as a bridge for forwarding data between VPC and the local data center.

Virtual extensible LAN

A simple network with MAC addresses encapsulated in UDP. It can create a virtual layer-2 subnet across multiple physical IP subnets.

Virtual Local Area Network

An emerging data exchange technology that logically divides Local Area Network (LAN) devices into segments to implement the virtual work group.

Virtual network layer

The network where the Elastic Compute Service (ECS) instances reside. It avoids the limits of physical network layer. User logic and transmitted contents at the virtual network layer are invisible to the physical network layer, and the physical network layer does not parse the transmitted information at the virtual network layer, which protects the user privacy. The virtual network layer is completely implemented by using softwares. The separation between data layer and control layer, path decision, and policy delivery are completed by the controller at the virtual network layer. Therefore, the Software Defined Network (SDN) is implemented.

Virtual Private Cloud (VPC)

The custom private network created based on Alibaba Cloud. Full logical isolation is achieved between VPCs. You can create and manage cloud product instances, such as Elastic Compute Service (ECS), Relational Database Service (RDS), and Object Storage Service (OSS), in your created VPCs.

VNC

A function in the Alibaba Cloud console, which is used to connect to the Elastic Compute Service (ECS) instances if common remote connection tools, such as PuTTY, Xshell, and SecureCRT, are unavailable.

VRouter

The Virtual Private Cloud (VPC) hub that can connect to all VSwitches in the VPC and serves as a gateway device that connects the VPC to other networks. It forwards the network traffic according to the configurations of route entries.

VServer group

A group of Elastic Compute Service (ECS) instances that process the frontend requests distributed by Server Load Balancer. Different listeners can use different VServer groups, which implements the request forwarding in the listening dimension.

VSwitch

A basic network device of Virtual Private Cloud (VPC), which can connect to different cloud product instances and must be specified when you create a cloud product instance in VPC.

W

Webshell

A segment of web codes (mostly the ASP and PHP codes) running on the server. Attackers perform malicious operations on the server by using this segment of codes to obtain sensitive technical information or gain the control over the server by penetration and privilege escalation.

WebTracking

A function that Log Service uses to support collecting logs from HTML, H5, iOS, and Android platforms and customizing dimensions and metrics.

Weighted Least Connections

A scheduling algorithm that Server Load Balancer uses to forward traffic. In addition to the weight set to each backend server, the actual loads (number of connections) of the backend server must be considered. If the weights of the backend servers are the same, a new request is sent to the server with the fewer connections.

Weighted Round Robin

A scheduling algorithm that Server Load Balancer uses to forward traffic. Backend servers with higher weights receive more requests than those with lower weights.

Workspace

- A container of Quick BI objects, such as data portals, spreadsheets, dashboards, data sources , and datasets. Multiple users can access the same workspace at the same time.
- An area where the computing resources and storage resources to be used by the algorithms are defined.